

Dispersion in Dispersion: Measuring Establishment-Level Differences in Productivity

by

Cindy Cunningham
U.S. Bureau of Labor Statistics

Lucia Foster
U.S. Census Bureau

Cheryl Grim
U.S. Census Bureau

John Haltiwanger
University of Maryland, NBER and IZA

Sabrina Wulff Pabilonia
U.S. Bureau of Labor Statistics and IZA

Jay Stewart
U.S. Bureau of Labor Statistics and IZA

Zoltan Wolf
New Light Technologies

This paper is a revised version of *Dispersion in Dispersion: Measuring Establishment-Level Differences in Productivity (CES 18-25RR)* from September, 2020. A copy of the original paper is available upon request.

CES 18-25RRR

August, 2022

The research program of the Center for Economic Studies (CES) produces a wide range of economic analyses to improve the statistical programs of the U.S. Census Bureau. Many of these analyses take the form of CES research papers. The papers have not undergone the review accorded Census Bureau publications and no endorsement should be inferred. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Republication in whole or part must be cleared with the authors.

To obtain information about the series, see www.census.gov/ces or contact Christopher Goetz, Editor, Discussion Papers, U.S. Census Bureau, Center for Economic Studies 5K028B, 4600 Silver Hill Road, Washington, DC 20233, CES.Working.Papers@census.gov.

Abstract

Abstract: We describe new experimental productivity dispersion statistics, Dispersion Statistics on Productivity (DiSP), jointly produced by the Bureau of Labor Statistics (BLS) and the Census Bureau, that complement the official BLS industry-level productivity statistics. The BLS has a long history of producing industry-level productivity statistics, which represent the average establishment-level productivity within industries when appropriately weighted. These statistics cannot, however, tell us about the variation in productivity levels across establishments within those industries. Dispersion in productivity across businesses can provide information about the nature of competition and frictions within sectors and the sources of rising wage inequality across businesses. DiSP data show enormous differences in productivity across establishments within industries in the manufacturing sector. We find substantial variation in dispersion across industries, increasing dispersion from 1997 to 2016, and countercyclical total factor productivity dispersion. We hope DiSP will enable further research into understanding productivity differences across industries and establishments and over time.

Keyword: manufacturing, reallocation, business cycles, productivity dispersion

JEL Classification: D24, E24, E32

* Cunningham: U.S. Bureau of Labor Statistics; Pabilonia and Stewart: U.S. Bureau of Labor Statistics and IZA; Foster and Grim: Center for Economic Studies, U.S. Census Bureau; Haltiwanger: University of Maryland, NBER and IZA; Wolf: New Light Technologies. John Haltiwanger was also a Schedule A part-time employee of the U.S. Census Bureau at the time of the writing of this paper. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or the Bureau of Labor Statistics. The Census Bureau has reviewed this data product for unauthorized disclosure of confidential information and has approved the disclosure avoidance practices applied to this release (Approval IDs: DRB-FY19-393, DRB-FY19-526, CBDRB-FY20-410). We thank John Earle, Kevin Fox, Bernd Görgiz, Bart Hobijn, Mark Roberts, Jon Samuels, Chad Syverson, the Federal Economic Statistics Advisory Committee, the BLS Technical Advisory Committee, and participants at the Conference on Research on Income and Wealth and the 2015 Federal Statistical Research Data Center Conference for their helpful comments.

1. Introduction

Productivity measures are critical for understanding economic performance in the U.S. economy. In this paper, we describe a new productivity data product, Dispersion Statistics on Productivity (DiSP), jointly developed and published by the U.S. Bureau of Labor Statistics (BLS) and the U.S. Census Bureau.¹ BLS produces the official labor and total factor productivity (TFP) growth statistics for major sectors and industries in the U.S. These statistics are constructed using aggregate industry-level data and can be thought of as changes in the first moment of establishment-level productivity (appropriately weighted). That is, these statistics show how productivity changes *on average* within sectors and industries, but they cannot provide insight into the *variation* in productivity levels across establishments within sectors or industries. Research has shown that changes in the dispersion of productivity across establishments in the same industry is related to changes in the growth of aggregate productivity (both economy-wide and at the industry level) through a variety of channels.

To fill this void, BLS and the Census Bureau initiated the Collaborative Micro-productivity Project (CMP) to develop and publish experimental statistics on within-industry productivity dispersion (i.e., second-moment measures of establishment-level productivity) and to produce restricted-use research datasets. The public-use statistics developed in this project, DiSP, were released for the first time in the fall of 2019. The first release covered all four-digit NAICS industries in the manufacturing sector and the years 1997–2016. The most recent release of DiSP in the fall of 2021 extended the coverage to 1987–2018.² Moving forward, the data will be updated annually. Restricted-use establishment-level data with

¹ DiSP data are available on both BLS and Census Bureau websites at: <https://www.bls.gov/productivity/tables/> and <https://www.census.gov/disp>. The DiSP is an experimental dataset and plans are underway to expand the data product in several ways as described in the paper.

² The additional earlier years in the 2021 release reflect the ongoing efforts to enhance the DiSP product.

micro-based estimates of productivity and its underlying components (e.g., output and input measures) are also available to qualified researchers on approved projects in secure Federal Statistical Research Data Centers (FSRDCs).³

Economic theory and recent empirical evidence suggest that the second moments of productivity are informative on several important dimensions. One of the most important findings in the literature on micro-level productivity is that large productivity differences across establishments exist even within narrowly defined industries.⁴ For example, using data from the 1977 Census of Manufactures (CM), Syverson (2004a) found that an establishment at the 90th percentile of the within-four-digit-SIC labor productivity distribution is on average about four times as productive as an establishment at the 10th percentile.

Syverson's findings generated considerable interest in the causes and consequences of this dispersion. Possible market explanations include the concavity of the profit function that prevents the most-productive business from taking over an industry, frictions in factor adjustments (such as costs of adjusting input factors), barriers to entry and exit, and distortions that inhibit the equalization of marginal products across businesses (such as the regulatory environment). Drivers of establishment-level productivity variation include differences in management skills, the quality of production factors, innovation, and investments in R&D.

Research has shown that the dispersion in establishment-level productivity varies across sectors, by geographic area, and over time. For example, Syverson (2004a, 2004b) shows that variation in dispersion across industries and geographic areas is related to product substitutability, market structure, and competition. Hsieh and Klenow (2009) argue that both cross-country variation and within-country variation in the dispersion in productivity are

³ For more information on the FSRDCs: <http://www.census.gov/fsrdc>. An earlier version of this dataset was analyzed in Foster et al. (2016a).

⁴ Syverson (2011) provides a survey of this literature.

related to distortions that inhibit productivity-enhancing reallocation. Asker et al. (2014) present evidence that the patterns of dispersion reflect dynamic factor adjustment frictions within sectors. The findings in Foster et al. (2016b) suggest that productivity differences across establishments may be generated by differences in efficiency levels, demand shocks, frictions/distortions, or all of the above. Foster et al. (2021a) and Cunningham et al. (2021) show that industries experiencing a surge in innovation exhibit a burst of firm entry, followed by an increase in productivity dispersion during an experimentation and shakeout phase, followed ultimately by an increase in industry-level productivity.

Establishment-level productivity differences are also correlated with important economic outcomes at the micro level, such as the survival and growth of establishments. There is an extensive literature on the connection between productivity, reallocation, and growth (Baily et al., 1992; Griliches and Regev, 1995; Foster et al., 2001; Diewert and Fox, 2010; Petrin et al., 2011; Foster et al., 2016a; Decker et al., 2020; Blackwood et al., 2021). These studies show that more productive businesses are more likely to survive and grow. These findings contribute to the perspective that reallocation—the process by which economic activity is allocated to its highest valued use—is an important contributor to aggregate productivity growth.

Productivity dispersion is also important for understanding rising wage inequality, which has been shown to be a between-firm phenomenon (Davis and Haltiwanger, 1991; Barth et al., 2016; Song et al., 2019; Haltiwanger and Spletzer, 2020). In addition, several studies have found that high-wage establishments are also more productive and that rising between-establishment dispersion in wages is closely associated with rising between-establishment dispersion in productivity (e.g., Dunne et al., 2004). Economic theories of search and matching provide the theoretical connection between productivity dispersion and wage dispersion (e.g., Burdett and Mortensen, 1998). Search and matching frictions create

quasi-rents for worker-firm matches that make it optimal for high-productivity firms to pay high wages.

Our results using the DiSP experimental data confirm earlier findings of sizeable differences in productivity across establishments within industries. To preview our results, we find that, on average, a manufacturing establishment at the 75th percentile of the within-industry labor productivity distribution is more than twice as productive as an establishment at the 25th percentile. If we instead focus on TFP, we find that an establishment at the 75th percentile is almost twice as productive as an establishment at the 25th percentile. Underlying these averages, we find substantial differences in dispersion across industries. For example, labor productivity dispersion in an industry at the 75th percentile of the dispersion distribution is about 1.4 times larger than the dispersion in an industry at the 25th percentile. The corresponding multiplier for TFP is 1.2. We also find that dispersion in within-industry productivity exhibits a positive time trend over our sample period (1997–2016) and that TFP dispersion is significantly countercyclical.

The experimental productivity dispersion statistics are intended to complement official BLS data, so it is crucial to understand the relationship between the dispersion of the productivity distribution derived from Census Bureau microdata and the statistics from BLS built from industry-level aggregates. Section 2 describes BLS productivity measures and productivity measures that we construct from Census microdata. Section 3 compares the two approaches to measuring inputs, output, and productivity for the manufacturing sector, and for four-digit NAICS manufacturing industries. We also compare these measures to data from the NBER-CES Manufacturing Industry Database and examine several data and measurement issues such as imputation and weighting of the microdata. In Section 4, we explore the variation in industry-level productivity dispersion measures across industries and over time. Section 5 summarizes our conclusions and describes plans for future work. In Appendix A,

we provide a table of acronyms used throughout the paper and their meanings for the ease of the reader.

2. Measuring Productivity

Because our primary goal is to create statistics that provide insights about productivity that complement the official BLS industry-level productivity measures, it is useful to first describe how BLS constructs its measures from published aggregates, and then compare it to measures that we construct by aggregating Census microdata.

2.1. BLS Industry-level Productivity

BLS publishes quarterly and annual measures of labor productivity growth for major sectors; annual measures of labor productivity for 199 three-digit and four-digit NAICS industries; and annual measures of TFP for major sectors, 18 three-digit NAICS manufacturing industries, 86 four-digit NAICS manufacturing industries, the air transportation industry, and the line-haul railroad industry. Productivity growth is measured as the difference between percentage changes in indexes of output and inputs (labor and, in the case of TFP, also capital and intermediate purchases). BLS does not publish industry productivity levels, although they are available on request.

BLS industry output is based on a sectoral concept, which measures the value of goods produced for sale outside the industry.⁵ For manufacturing industries, BLS uses published Annual Survey of Manufactures (ASM) and Census of Manufactures (CM) data on the total value of shipments, which it adjusts to remove intrasectoral transactions and resales and to account for changes in finished goods and work-in-process inventories.⁶ This adjusted

⁵ Sectoral output is less than gross output, but greater than value-added output. In the most detailed industries, sectoral and gross output are the same or very close. However, going from very detailed industries to more aggregated industries, sectoral output moves closer to value-added output. In the limit, at the aggregate level, sectoral output is the same as value-added output, except for imported intermediate inputs. For more information on the importance of using sectoral output, see Kovarik and Varghese (2019).

⁶ See <https://www.census.gov/programs-surveys/asm.html> and <https://www.census.gov/programs-surveys/economic-census.html>.

nominal output measure is then distributed to detailed categories of products and services using the mix of annual wherever-made product shipments from the ASM. Nominal output in each product category is deflated using the appropriate detailed producer price index from the BLS prices program. These real output measures are then Tornqvist-aggregated into industry output indexes. Self-employment revenues for manufacturing firms, which come from Internal Revenue Service data, are also added to these output measures.

BLS measures labor input as the total annual hours worked by all persons in an industry. This measure is constructed by combining data from three BLS surveys: the Current Employment Statistics (CES) survey, the Current Population Survey (CPS), and the National Compensation Survey (NCS). The CES provides detailed information on the employment and average weekly hours *paid* for production and non-supervisory employees (henceforth referred to as production workers).⁷ The NCS data are used to adjust CES average weekly hours from an hours-paid to an hours-worked basis by removing paid vacation accrued and sick leave taken.⁸ To estimate nonproduction worker average weekly hours, BLS uses data from the CPS to calculate a ratio of nonproduction to production worker average weekly hours worked, which is then multiplied by the adjusted CES production worker hours (worked). Total nonproduction worker hours are estimated as:

$$TH_{NP} = \text{Emp}_{NP}^{\text{CES}} \times AWH_P^{\text{CES}} \times hwhp_P^{\text{NCS}} \times \frac{AWH_{NP}^{\text{CPS}}}{AWH_P^{\text{CPS}}} \times 52 \quad (1)$$

where $\text{Emp}_{NP}^{\text{CES}}$ is nonproduction worker employment from CES, AWH_P^{CES} is production worker average weekly hours paid from CES, $hwhp_P^{\text{NCS}}$ is the hours-worked-to-hours-paid ratio from NCS, and $\left(\frac{AWH_{NP}^{\text{CPS}}}{AWH_P^{\text{CPS}}}\right)$ is the CPS nonproduction/production average weekly hours ratio. CPS data are also used to construct hours worked by self-employed and unpaid family

⁷ Workers in goods-producing industries are referred to as being production or non-production employees and in the service-providing industries as nonsupervisory or supervisory employees.

⁸ Note that this adjustment does not account for off-the-clock hours.

workers (Eldridge et al., 2004).

For TFP at the four-digit NAICS level, capital input is based on the flow of services from the productive stock of capital. BLS investment data for industries combine expenditures on structures and equipment from the ASM with data on investment in different assets by industry from BEA and the Annual Capital Expenditures Survey (ACES). Using a perpetual inventory method, BLS then computes industry-asset level capital stocks from these investment flows. In BLS official TFP measures, these stocks are converted to capital services using industry-asset specific rental prices and then aggregated to the industry level.

For intermediate purchases inputs, BLS combines quantities of materials, purchased business services, fuels, and electricity consumed by each industry. The nominal values of materials, fuels, and electricity are from the CM and ASM, while the values of purchased business services are estimated from BEA and Census Bureau data.

2.2. Establishment-level Productivity using Census Data

To measure establishment-level labor productivity, we combine establishment-level information from three Census Bureau restricted-use microdata files with public-use industry-level data from BLS. Given that one goal of our research is to shed light on BLS industry productivity statistics, we try to match BLS concepts and measures as closely as possible.

Our establishment-level microdata come from the CM, the ASM, and the Longitudinal Business Database (LBD). The CM is collected every 5 years in years ending in “2” and “7”. Data are collected from all manufacturing establishments except those that are very small. For these very small non-mail establishments, the Census Bureau uses information from administrative records. The ASM sample is a 5-year panel of manufacturing establishments, updated every year for births, and data are collected annually. ASM panels begin in years ending in “4” and “9”, and the probability of selection into the ASM sample is a function of both industry and size (employment or the value of shipments). Like the CM,

the ASM does not collect data from very small establishments but accounts for them using administrative information. In CM years, ASM data are collected as part of the CM data collection, but for this analysis and the public-use statistics, we use only the ASM establishments.⁹ Data are imputed for establishments that do not respond or that fail to report some data elements (item non-response); we discuss this further in Section 2.3. The LBD is a longitudinally linked version of the Census Bureau’s Business Register that covers the non-agricultural employer universe of business establishments (see Jarmin and Miranda, 2002 and Chow et al., 2021). The LBD provides us with both high-quality longitudinal links and information on the universe of manufacturing establishments, which we use to construct the inverse propensity score weights (IPW) that we use in our productivity calculations.

Ideally, we want to construct an output measure that exactly matches the BLS measure. We start by using Census microdata to replicate the value of shipments as closely as possible. Specifically, we calculate establishment-level real output as deflated revenues, adjusted for resales and changes in inventories.¹⁰ However, we cannot replicate the BLS sectoral output concept because the ASM does not collect the information needed to calculate intra-sectoral transactions. Instead, we add the value of intra-sectoral transactions back into BLS output measures to make the two measures comparable. Thus, we measure establishment-level output as:

$$Q_{et} = (TVS_{et} + DF_{et} + DW_{et} - CR_{et})/PISHIP_{it} \quad (2)$$

where TVS = total value of shipments, $DF_{et} = FIE_{et} - FIB_{et}$ and $DW_{et} = WIE_{et} - WIB_{et}$ are the changes in finished-goods and work-in-process inventories, respectively (FIB , FIE = beginning-of-year and end-of-year finished goods inventories, and WIB , WIE = beginning-of-year and end-of-year work-in-process inventories), CR = cost of resales, $PISHIP$ = deflator

⁹ The microdata made available in the FSRDCs contains productivity measures for all CM establishments for which productivity calculation is possible.

¹⁰ In practice, subtracting resales does not make much difference because they are only a small fraction of revenue.

for the value of shipments, and the i , e , and t subscripts index industries, establishments, and years, respectively.¹¹

We measure labor input as total hours worked. For each establishment, the ASM collects the total number of employees, the number of production workers, and the total number of hours worked by production workers. We calculate total annual hours worked by summing ASM production worker hours and an estimate of nonproduction worker hours, which we calculate using the same methodology as BLS (equation (1)) but substituting ASM data for CES data:¹²

$$TH_{et} = PH_{et} + \left((TE_{et} - PW_{et}) \times \frac{PH_{et}}{PW_{et}} \times \left(\frac{AWH_{NP}^{CPS}}{AWH_P^{CPS}} \right)_{it} \right) \quad (3)$$

where PH = production worker hours, PW = the number of production workers, TE = total

employment, and $\frac{AWH_{NP}^{CPS}}{AWH_P^{CPS}}$ = the CPS non-production/production average weekly hours ratio

for the four-digit NAICS industry. We calculate establishment-level log labor productivity as:

$$\ln LP_{et} = \ln Q_{et} - \ln TH_{et}. \quad (4)$$

Establishment-level TFP in logs is measured as:

$$\ln TFP_{et} = \ln Q_{et} - \alpha_K \ln K_{et} - \alpha_L \ln TH_{et} - \alpha_M \ln M_{et} - \alpha_E \ln E_{et}, \quad (5)$$

where Q and TH are real output and total hours as defined in equations (2) and (3), K denotes

real productive capital stock, and M and E denote deflated values of expenditures on

materials and energy, respectively.¹³ The productive capital stock is constructed using the

perpetual inventory method for equipment and structures separately.¹⁴ The value of M is

¹¹ The value of shipments are deflated at the six-digit NAICS industry level using deflators from the NBER-CES dataset. Values for 2012–2016 are imputed using price indexes from BEA GDP-by-Industry data. In comparisons to BLS output data conducted in this paper, BLS industry implicit output price deflators (as described in 2.1) are used in place of the shipments deflator.

¹² Note that it is not necessary to use the hours-worked-to-hours-paid ratio because establishments are requested to report hours worked.

¹³ BLS constructs aggregate TFP using capital services rather than productive stock.

¹⁴ We do not include rented capital due to its irregular collection in the ASM. Pre-1986 and post-2006, this information is collected annually on the ASM. In the intervening years, this information was only collected in the Economic Census. Exploratory analysis for the years when this is available shows rented capital is small and

calculated as the deflated sum of the cost of materials, the cost of resales, and the cost of contract work done for the establishment by others.¹⁵ The nominal value of E is the sum of the cost of electricity and the cost of fuels. For the CMP data ultimately included in the DiSP data, the two expenditures are deflated using the appropriate deflators from the NBER-CES database. Equation (5) follows directly from a Cobb-Douglas production function, a common assumption in the literature.¹⁶ We measure the factor elasticities, α_K , α_L , α_E , and α_M , using the share of expenditures of the corresponding input in total cost in each six-digit NAICS industry.¹⁷ For the BLS data comparisons, we use published BLS industry cost shares; for the DiSP statistics, we use shares calculated from CMP data.

DiSP does not include value-added-based productivity dispersion for conceptual reasons. While there is a market for final demand, a market for value added does not exist. This consideration is irrelevant for the overall economy because value added equals the value of final demand at that level of aggregation. However, the relationship between value added and final demand is only an approximation at lower levels of aggregation. In order to make inferences about final goods produced by an establishment or industry, not only are the output and input values necessary but also the relevant input-output linkages.¹⁸ A related issue is that

does not make much difference to establishment-level capital measures. We plan on exploring this further in future research.

¹⁵ Published ASM measures of value added are based on the difference between the value of output and a composite-operating-expenses measure inclusive of materials and energy expenditures. We break out the cost of materials and energy separately in our TFP measure. The inclusion of contract work implies that some aspects of purchased services are included in the materials expenditures. However, since 2006, the ASM has included questions on other operating expenses, including leased employees and additional purchased services not included in the cost of contract work. We are actively exploring the inclusion of those operating expenses for a supplemental TFP dispersion measure commencing in 2006. Challenges for inclusion of these variables are the short time series, item non-response rates, and the treatment of establishments of single-unit versus multi-unit firms. Establishments of the latter are less likely to have such additional operating expenses because the headquarters establishment of the parent firm may be providing and/or purchasing those services.

¹⁶ It can be shown that the Cobb-Douglas production function is a first-order approximation of any CES production function. A popular alternative, the translog, is a second-order approximation; see Caves et al. (1982). While the translog is more flexible, it requires estimates of second derivatives and is therefore a less parsimonious specification.

¹⁷ See web appendix B in Foster et al. (2016a) for more details. Procedures for extending the output and input price deflators are described in this appendix.

¹⁸ An additional consideration follows when using the logarithmic transformation to stabilize the variance of

the existence of a value-added production function at the establishment-level requires very strong functional form assumptions that are likely violated (Basu and Fernald, 1997).

2.3 Missing Data and Imputation

As noted earlier, the ASM microdata are subject to item non-response, and these missing values are imputed by the Census Bureau. The Census Bureau’s imputation methods are designed to yield accurate published aggregates but do not necessarily preserve the distribution or adequately reflect the variability of the underlying microdata. There is evidence that certain imputation methods may affect microdata analyses. However, there are techniques available to mitigate the effects of imputation on dispersion measures. For example, White et al. (2018) analyze dispersion statistics using classification and regression-tree methods. Blackwood et al. (2021) follow a different approach and address imputation by dropping observations with imputed data and reweighting the remaining observations. The results from these studies suggest that imputation yields lower measured dispersion relative to the case when imputation is corrected for. In this paper, we use the entire set of observations in the sample and leave further analysis of these issues for future work.

3. Comparing Micro-Aggregated Data to Published Industry Data

In this section, we compare our micro-aggregated estimates to the official data published by BLS, covering the 1997–2016 period. Based on earlier work comparing similar business data across the two government agencies, we expect that there will be some systematic differences between these measures (Elvery et al., 2006). Even though differences in the levels of the micro-aggregated and published first moments do not directly affect our conclusions about dispersion (because we control for industry-year effects), it is useful to determine how far apart the two sets of estimates are. If the first moments are close, then it is

empirical distributions. This transformation truncates distributions at zero, which leads to biased inference if the probability of negative value added is correlated with establishment-level characteristics. Negative value added may be plausible, for example, in intermediate-intensive industries during periods of recession or crises (Cao et al., 2021).

more reasonable to think of the micro-based second moments as measuring variation around the published first moments. As an additional check, we compare these two sets of estimates to estimates from the NBER-CES database. The NBER-CES estimates can be thought of as equivalent to the official published ASM and CM statistics upon which they are based and are used here for comparison only.¹⁹ We start by comparing input and output measures, and then we compare productivity measures.

3.1. Input and Output Measures

Figure 1 shows the total number of employees in the manufacturing sector from each series. The first thing to note is that employment levels based on ASM microdata (using ASM sample weights) are significantly lower than the published ASM and BLS estimates because they exclude the “non-mail” stratum—small establishments that are not sampled by the ASM. The published ASM series includes adjustments for the non-mail stratum and is much closer to the BLS estimates.

To account for these small non-mail establishments, we construct an alternative set of weights. As noted, the weighted sample total calculated from the ASM (using ASM sample weights) is by design not equal to the published total because there are additional adjustments in the latter for the non-mail cases. Fortunately, there are some very small establishments in the ASM sample each year that are below the thresholds for non-mail cases. This occurs because, among the smaller establishments that were selected for the ASM (that is, establishments with employment above the threshold), some had fallen below the size threshold by the time that they provided data. This implies that there is coverage for all business sizes in the ASM sample (e.g., there are ASM establishments in any given year and industry with 1–4 employees, even though this is typically below the ASM sample threshold).

¹⁹ For more information on the NBER-CES Manufacturing Industry Database, see <http://www.nber.org/nberces/>. The NBER-CES series was last updated through 2011 as of September 15, 2020.

To create the alternative set of weights, we use the LBD. Specifically, we define the manufacturing universe using the LBD and use LBD data to estimate the probability that an establishment is included in the ASM sample. We then use these probabilities to construct IPW. See Appendix B for a full discussion of the weighting procedures. This procedure increases the weights assigned to the “non-mail” establishments so that the propensity-score-weighted employment totals are consistent with the LBD.²⁰ Moreover, as seen in Figure 1, the micro-aggregated employment series using IPW yields totals that align with the published BLS and ASM more closely than those using ASM weights.²¹ This aligns with our objective in using IPW, as these weights correct for the contribution of the “non-mail” establishments in a manner that the ASM weights are not designed to address.

Next, we compare total manufacturing output and input growth across the BLS, CMP, and NBER series (Figure 2).²² For all three series, we use the BLS price deflators for these comparisons, apart from the capital series, because BLS does not produce separate deflators for equipment and structures. To make it comparable to what we can construct from the ASM, we also adjust the BLS output series by adding the value of intrasectoral transactions back into the series. The BLS and NBER output series track each other closely, while the CMP series deviates from the other two series in some years but exhibits the same pattern of growth rates (see Panel (a)). Panel (b) of Figure 2 compares hours growth rates, which exhibit similar dynamics across the series, except during the 2005–2007 period, when the NBER series diverges. Panel (c) of Figure 2 shows capital stock growth rates. All three series exhibit a slight downward trend in the late 1990s and are essentially flat starting in the early 2000s.

²⁰ In addition, unreported results suggest that IPW do a good job matching the industry/year-specific size and age distributions of the LBD.

²¹ We explored the possibility of benchmarking CMP employment (based on the manufacturing universe in the Census Business Register) to BLS employment (based on the manufacturing universe in the BLS Quarterly Census of Employment and Wages). While this benchmarking would improve the correlation between the BLS and CMP labor (employment and hours) measures, it decreases the correlation between BLS and CMP output and the other input measures, which are based on the manufacturing universe in the Census Business Register.

²² See Appendix A Table A2 for further details about the construction of these series.

Panels (d) and (e) of Figure 2 show that for energy and material inputs, the three series track very closely.²³

Table 1 shows correlations between the three data sources for inputs and output for the total manufacturing sector. The correlations in the top panel are based on total manufacturing aggregate time series, while the bottom panel shows the average of the within-industry correlations for four-digit NAICS industries, calculated over the 19 years of the sample. The top panel of the table shows that hours, energy, materials, and output, both in levels and in growth rates, are highly correlated across the data series (the correlations range from 0.88 to 0.99).²⁴ Average industry-level correlations, shown in the bottom panel, are lower than for total manufacturing, but they are still reasonably high for these variables, both in levels and in growth rates.

The correlation between the capital series is significantly lower than the correlations described above. There are several plausible explanations. First, there is a fundamental difference in the underlying data. BLS investment data for four-digit industries combine expenditures on structures and equipment from the ASM with data on investments in different assets by industry from BEA and the Annual Capital Expenditures Survey (ACES).²⁵ Using a perpetual inventory method, BLS then computes industry capital stocks from these investment flows. In contrast, our approach takes investment flows directly from the establishment and uses these flows with the perpetual inventory method at the establishment level to generate capital stocks. Second, the BLS investment series covers a longer period than the micro-aggregated series. This is significant because initial capital stocks are difficult to measure. The BLS capital stock is built up from investment flows that stretch back to 1958 (and longer for some assets). By 1997, when our data series starts, the

²³ The cost of purchased services and resales are not included in the materials comparisons.

²⁴ The level correlations being high reflects trends, implying appropriate caution in interpretation.

²⁵ See Becker et al. (2006) for more discussion of the relationship between top-down methods used by BEA and bottom-up methods like those we use with the establishment-level data for measuring capital stocks and flows.

BLS capital stock reflects mainly investment, and the impact of any mismeasurement of the initial capital stock is minimal. In contrast, the sample rotation in the ASM implies that we need to estimate initial capital stocks for establishments that newly enter the sample rotation. For CMP, we initialize new establishments using their book value, and the earliest book value and investment data that we use dates to 1972. Despite differences in data sources and methodologies, we can conclude that the micro-aggregated data are largely consistent with published aggregate data.

3.2. Productivity Growth

We calculate productivity growth as the change in log productivity where productivity is measured as either output per hour or TFP.²⁶ Figure 3(a) shows that output-per-hour growth rates for the manufacturing sector are broadly similar, but with some greater discrepancies in various subperiods (e.g., 2003–2009). These differences can be attributed to the differences in data sources and methodologies, some of which were illustrated in Figure 2. Despite underlying differences, TFP growth shows remarkable similarity across these data sources, see Figure 3(b). Table 2 echoes these findings: the correlations between the series constructed using different data sources are highest for TFP growth.

This comparison of inputs, output, and productivity serves as an important backdrop to our new experimental statistics on within-industry dispersion. Although there are some differences between the BLS aggregates and the micro-aggregated series, they are similar enough to allow us to make meaningful inferences about the relationship between within-industry dispersion and BLS published estimates of industry productivity growth.

4. Productivity Dispersion

For our analysis of productivity dispersion, we focus on (log) levels rather than

²⁶ The official BLS productivity series are calculated using percentage changes in the index, and thus the BLS series that we refer to here differ slightly from the published series. Additionally, the official total manufacturing productivity series is published by the BLS Division of Major Sector Productivity, whereas the data here are aggregated from industry data provided by the BLS Division of Industry Productivity Statistics.

growth rates. Because we are interested in comparing within-industry dispersion of productivity across industries and over time, it is necessary to account for industry differences in average productivity. To do this, we calculate establishment-level productivity as the deviation from average productivity in that establishment's four-digit industry in each year.²⁷ The interpretation of normalized productivity levels is intuitive: they tell us how far above or below the mean the establishment sits in the productivity distribution.

We use the interquartile range (IQR) as our primary measure of dispersion, because it is intuitive and easy to interpret. The IQR shows how much more productive an establishment at the 75th percentile of the productivity distribution is than an establishment at the 25th percentile of the productivity distribution. The standard deviation may seem like an obvious alternative to the IQR; however, it is known to be more sensitive to outliers than quantile-based dispersion measures. We also report the 90–10 differential as well as the 10–1 and 99–90 differentials.

Table 3 shows the descriptive statistics for distributions of the dispersion measures.²⁸ The first entry reported in the table (0.898) indicates that in the average industry and year, an establishment at the 75th percentile of the distribution is about ($e^{0.898} \approx$) 2.45 times as productive as an establishment at the 25th percentile. An establishment at the 90th percentile is about 5.9 times as productive as one at the 10th percentile. Average dispersion in TFP is lower.²⁹ However, an establishment at the 75th percentile is still about 1.7 times as productive as an establishment at the 25th percentile using TFP measures. These differences imply substantial differences in a core measure of business performance at the establishment-level within narrowly defined industries.

²⁷ These are weighted averages using IPW where establishment-level productivity is expressed as a deviation from average productivity in that establishment's 4-digit NAICS industry.

²⁸ We present standard deviations in Table 3 for completeness but do not discuss these results.

²⁹ The range for output per hour is somewhat larger than what was found by Syverson (2004a)—he found a multiplier of about 1.9—but our findings are generally in line with his results.

Many factors may underlie the observed dispersion in measured productivity across establishments in the same industry.³⁰ We define a “wedge” as any mechanism that prevents the equalization of marginal revenue products across producers. Because the measures of productivity dispersion reported here are revenue-based measures, the presence of widespread dispersion is consistent with the presence of one or more types of wedges. One type is adjustment frictions that inhibit businesses from adjusting their scale of operations and specific inputs to changing economic conditions. These adjustment frictions may be related to the costs of adopting new technologies or business practices; thus, dispersion in an industry may reflect the gap between the frontier establishments and other producers. Additional types of wedges are market distortions such as differences in markups across producers or financial constraints in the same industry. Complicating matters is that in the presence of wedges that are correlated with fundamentals, the variation in the dispersion will also reflect differences in business fundamentals such as technical efficiency and product appeal across businesses (Blackwood et al., 2021). For example, an increase in the dispersion in product appeal across producers in the presence of adjustment frictions will yield an increase in the dispersion of revenue productivity across producers (even if the adjustment frictions remain constant). In a similar fashion, dispersion may reflect unmeasured inputs. These could include production methods, management practices, and the mix of worker types (labor quality). Rising revenue productivity dispersion might also reflect rising dispersion in firm-level markups (De Loecker et al., 2020; Foster et al., 2021b).

It is well beyond the scope of the current paper to determine the relative importance of each of these alternative factors. Instead, the aim here is to describe the DiSP data product and point to its potential for investigating these alternative determinants of dispersion. One

³⁰ See Syverson (2011) for more detailed discussion of these issues and Blackwood et al. (2021) for a discussion more closely related to this new data product.

strength of the new data product is that dispersion measures are provided at a detailed level of aggregation by year. Figure 4a summarizes how within-industry dispersion in output per hour—measured as the IQR—varies across industries and over time. The mean and median IQR are close to each other, but the large differences between the 25th and 75th percentiles show that there is substantial variation in the IQR across industries. For example, in 2002, the productivity difference between an establishment at the 75th and 25th percentiles is about 100 log points in an industry at the 75th percentile of the IQR distribution, while this difference is approximately 70 log points in an industry at the 25th percentile of the IQR distribution. We find similar differences across the industry distribution when looking at TFP dispersion, with IQRs of approximately 60 and 40 log points at the 75th and 25th percentiles, respectively (see Figure 4b).

The differences in the IQRs suggest that there are factors such as those discussed above that generate “dispersion in dispersion,” including differences in shocks, adjustment costs, distortions, technology, and distributions of capital intensities. In addition, dispersion is rising during the period under investigation, more so for TFP than for labor productivity. The rising trend suggests that wedges, and the dispersion in business fundamentals underlying the observed dispersion, are changing in systematic ways over time.³¹ We can also see from Figure 4 that although the volatility of (the mean) dispersion is nontrivial, it is dwarfed by the variation across industries.

Table 4 shows that there is substantial movement in the ranking of industries in terms of their dispersion. The diagonal elements of these matrices are less than one, indicating that the probability that the IQR of an industry remains in the same quintile

³¹ There is an ongoing debate about the source of the rising dispersion in revenue productivity measures. See, e.g., Bils et al. (2021), Blackwood et al. (2021), Decker et al. (2020), and von Brasch et al. (2020). We do not seek to address that debate here directly but note that Decker et al. (2020) find that rising (revenue-based) labor productivity dispersion in manufacturing is present in both the ASM data used here and in administrative data from the Business Register. This finding suggests that an increase in measurement error is not driving the rising dispersion.

between two periods is less than one. Conversely, off-diagonal elements are generally non-zero. For example, the second entry of the first row indicates that there is a 24 percent chance that an industry in the first quintile of output per hour IQR last year is in the second quintile this year. Similarly, the fourth entry in the fifth row indicates that there is an 18 percent chance that an industry in the fifth quintile of output per hour IQR last year is in the fourth quintile this year. These findings illustrate that not only is there dispersion in dispersion, but the IQR rank of industries varies over time.

For the rest of this section, we consider a few extensions to our analysis to illustrate further the nature of the dispersion. We first examine how our results change when we weight establishments using activity weights. Activity weights are generated by multiplying our IPW by an activity measure such as hours shares (the share of an establishment's hours of the total hours in its industry) for labor productivity and composite input shares for TFP. Activity weighting paints a potentially different picture of dispersion because there may be differences between the dispersion of different size groups. Our second extension is to examine the tails of the productivity distribution. There has been great interest in the finding that a substantial portion of wage inequality is driven by the upper tail of the distribution and by increasing between-establishment wage differentials. Investigating the upper tail of the productivity distribution is analogously interesting, as theory and evidence show that the productivity and earnings distributions are related. Third, we explore the common secular trends and cyclicalities in the dispersion measures. Fourth, we highlight that most of the within-industry dispersion in productivity is not accounted for by establishment observables like business size, age, or location. Finally, we provide additional information about the public domain DiSP data product.

4.1. Activity-Weighted Dispersion Measures

Figure 5 replicates the dispersion measures in Figure 4 but using activity weights,

which we will refer to as “activity-weighted” distributions. We see rising dispersion over time in both Figures 4 and 5. The main difference is that activity-weighted dispersion is smaller and exhibits less year-to-year variation than dispersion weighted by IPW alone. Because activity weights give more weight to larger establishments, comparison of Figures 4 and 5 implies that there is less dispersion among larger establishments. Or put differently, dispersion in productivity across hours or composited inputs is lower than dispersion across establishments. Activity-weighted dispersion also exhibits a rising trend that, but it is less pronounced than the increase in trend dispersion without activity weighting.

4.2. Dispersion in the Tails

Turning to the tails of the productivity distributions, a distinctive feature of the within-industry productivity distribution is that mean and median dispersion in the right tail (the 99–90 difference—see Figure 6) are about the same order of magnitude as the mean and median dispersion of the IQR (see Figure 4). This is striking given that each tail covers only one-fifth as many establishments as the IQR. Comparing Figure 6a to Figure 4a, output per hour dispersion in the 99–90 difference is only slightly smaller than the IQR. TFP 99–90 differences are even larger than the IQR. In addition, differences in the right tail of the TFP distribution are rising faster: the mean shows that dispersion in the right tail rose by about 40 log points between 1997 and 2016, while the IQR rose by only 10 log points over the same period (compare Figures 6b and 4b). The activity-weighted dispersion measures (Figure 7) generally show similar patterns but a smaller absolute increase. In addition, there is substantial dispersion in dispersion in the right tail.

The fact that there are systematic differences in dispersion between different parts of the productivity distribution—for example, dispersion among the most-productive establishments is generally higher and is rising faster over time than dispersion among their less productive competitors—is relevant for important questions about aggregate productivity

growth. As mentioned in the introduction, earlier studies established a connection between reallocation and productivity growth. In a well-functioning market economy, production inputs flow from less-productive to more-productive businesses. In other words, gains in aggregate productivity due to resource reallocation are possible only when establishments have different productivity levels, which implies dispersion is relevant in this context because it can be thought of as an indicator of the potential gains from reallocation.³² Quantifying the contribution of different establishment types is an important empirical question and exploring their relative importance remains a promising area for the empirical productivity literature.

In contrast to the right tail, the left tail (the 10–1 ratio) exhibits lower dispersion relative to the IQR. Mean output-per-hour 10–1 differences are 20–30 log-points smaller than the IQR, though they exhibit similar volatility (see Figure 8a). Mean TFP differences in the left tail are similarly smaller than the IQR and exhibit no positive trend (see Figure 8b).³³ The activity-weighted dispersion measures tell a similar story (Figure 9).

These findings highlight the importance of looking at the entire productivity distribution. The IQR is a convenient measure that covers half of the distribution. However, there is as much dispersion in the upper and lower tails as there is in the center of the distribution. We also see that weighting matters: accounting for size tends to reduce both productivity dispersion and its volatility.

4.3. Common Secular Trends and Cyclicalities in Dispersion

As is evident in Figures 4 and 5, the moments of industry-level dispersion exhibit considerable variation over time with an apparent upward trend in the first moment. Table 5 quantifies common trends and cyclicalities in the IQR measures of labor productivity and TFP

³² Dispersion across establishments may reflect the frictions impeding efficient reallocation. Mitigating such frictions can improve productivity.

³³ The spike in mean dispersion in 1998 is due to transitory changes in the following four-digit NAICS industries: 3341, 3342, 3344, 3345, and 3351. In these industries, the least-productive establishments shifted to the left in 1998 and then back to the right in 1999. The significant changes in production technologies in these industries (factoryless manufacturing and offshoring) may explain these transitory dynamics in these years.

on both an activity-weighted and non-activity-weighted basis. These regression results are based on the public-domain DiSP panel data at the four-digit-NAICS-by-year level. All reported results include (unreported) industry fixed effects. A statistically significant positive trend is present for all measures considered. Cyclicalities are captured by the change in the national unemployment rate (Bureau of Labor Statistics, 1996–2016). Periods of increases in the unemployment rate correspond closely to NBER-defined recessions (Foster et al., 2016a). Dispersion in TFP is significantly countercyclical in the DiSP data, which is consistent with Kehrig (2015). Dispersion in labor productivity has a less consistent cyclical pattern. It is significantly procyclical when using the non-activity-weighted dispersion measure, but acyclical using the activity-weighted dispersion measure.

Many factors may underlie the countercyclical nature of productivity dispersion. It might reflect cyclicalities of second moment shocks that has been interpreted as countercyclical uncertainty shocks by Bloom (2009). It might also reflect countercyclical increases in frictions and distortions (see, e.g., Blackwood et al. (2021)) and in turn imply countercyclical increases in allocative inefficiency.³⁴

4.4. Establishment and Firm Characteristics and Within-Industry Dispersion

Many factors may underlie the substantial dispersion in productivity across establishments within the same industry as well as the variation between dispersion measures over time. To provide more guidance on the potential driving forces, we examine the relationship between productivity and observable establishment characteristics. For this purpose, we examine differences across space (by state), establishment size, and age.

Table 6 shows the R-squared and p-values from F-tests for regressions of establishment-level productivity on observable establishment characteristics, such as

³⁴ Diewert and Fox (2018) provide independent methodology and evidence that allocative inefficiency increases in recessions.

geography (state), size, and age classes. We focus on accounting for the variation in productivity within (four-digit NAICS) industry cells. While we find that geography, size, and age have statistically significant relationships with productivity variation across establishments in the same year and industry, these characteristics account for only a fraction of the observed differences in the productivity levels across establishments. It is this type of finding that highlights the interpretation in the research literature that there is enormous idiosyncratic variation in measured productivity across establishments. As we have discussed above, such idiosyncratic variation may stem from many factors. Further research is necessary to understand the contribution of different factors to this idiosyncratic dispersion. Table 6 highlights that much of the idiosyncratic dispersion in dispersion is industry specific and that the annual, industry-level statistics potentially provide much scope for investigating the sources of this dispersion.

4.5 Description of the Dispersion Statistics on Productivity (DiSP) Data Product

The new data product, Dispersion Statistics on Productivity (DiSP), contains a balanced panel of productivity statistics summarizing the within-industry distributions of output per hour and TFP.³⁵ Dispersion statistics include standard deviations, interquartile ranges, and interdecile ranges of the within-industry distributions of establishment-level productivity. All data moments are weighted using IPW so that they are representative of the universe of establishments (see Section 2.2 and Appendix B.2). In addition, the dataset includes activity-weighted versions of dispersion measures. We plan to include the 99–90 and 10–1 ranges in future releases, given the interesting patterns in the right and left tails highlighted above.

³⁵ The timeliness of the DiSP data depends on the release of establishment- and firm-level information. Our goal is to provide annual updates. In non-Census years, the ASM is available in the fall of the following year, while the LBD becomes available in spring of the year after that. In Census years, microdata become available later. The productivity dataset can be created approximately 2–3 months after the underlying microdata become available.

The data product is useful for analyzing the relationships between productivity dynamics at the establishment-level, industry-level, and for the entire manufacturing sector. As discussed above, many factors may underlie the cross-industry and time-series variation in dispersion. We expect that this new data product will facilitate our understanding of the connection between micro- and macro-level productivity. A key benefit of making these data available will be to allow researchers without access to the confidential microdata to explore the various possible causes—and effects—of the differences in within-industry dispersion across industries and over time.

5. The Relationship between Productivity Growth and Dispersion

The large and rising productivity dispersion discussed in the previous section could be due to one or more underlying mechanisms. For example, if dispersion results from innovative activity and experimentation that increases heterogeneity, then dispersion is a positive sign because innovative industries are likely to exhibit growth after a shakeout period (Gort and Klepper, 1982; Foster et al., 2021a; Cunningham et al., 2021). As mentioned earlier, dispersion may also be due to frictions and distortions (Hsieh and Klenow, 2009) that prevent the flow of resources from less productive to more productive businesses. In this case, dispersion has negative consequences for growth.

The new data product is well suited for analyzing the link between productivity growth and dispersion. To illustrate this, we regress BLS estimates of industry-level productivity growth on contemporaneous and lagged values of industry-level dispersion growth measures and other control variables (Table 7). We allow coefficients to differ between high-tech and non-tech industries and before/after the Great Recession. We also control for period effects. The pre-recession coefficients in column 2 show that an increase in LP dispersion is not associated with statistically significant productivity growth in non-tech industries but is associated with significant positive growth in high-tech industries. The

coefficient on lagged dispersion growth suggests that an increase in dispersion is followed by lower growth in non-tech industries but additional growth in high-tech industries; both correlations are statistically significant. The relationship between TFP dispersion and growth in high-tech industries in column 4 is not inconsistent with these findings, although the correlations are lower in absolute value and are estimated with less precision. In the post-recession period, the correlation structure is different: the non-tech industry correlations for LP are both positive and statistically significant while the high-tech industry correlations are not statistically significant (column 6). The signs of the high-tech correlations change for TFP (column 8). While pre-recession data show that rising productivity dispersion was followed by periods of positive TFP growth in high-tech industries, dispersion appears to have negative growth implications after the Great Recession.³⁶

The estimates in Table 7 are associations and do not imply causality. These patterns may be consistent with one or more of the above-mentioned mechanisms. For example, whether the positive link between growth and dispersion in high-tech industries is a sign of increased innovation cannot be established without measures of innovative activity or analyzing the role of entry.³⁷ Similarly, whether the negative relationship between growth and dispersion after the Great Recession reflects misallocation of resources across establishments cannot be assessed without analyzing the sources of dispersion in marginal revenue products. Analogously, for any explanation of the negative relationship using cyclical mechanisms or changes therein, see for example Kehrig (2015), a joint analysis of input prices and dispersion would be required. Nevertheless, the results in this section show systematic patterns between the first and second moments of the within-industry productivity distribution. The four-digit-

³⁶ The different industry and time patterns of these correlations are illustrated in Figure 10, where the growth in BLS TFP and within-industry dispersion are plotted for two high-tech industries (semiconductors and other electronic component, computer and peripheral equipment) and two non-tech industries (motor vehicles parts, fabric mills).

³⁷ Foster et al. (2021a) and Cunningham et al. (2021) analyze the relationship between dispersion and entry in U.S. industries.

NAICS-by-year DiSP data covering several decades offers considerable scope for investigating these and other hypotheses.³⁸

6. Concluding Remarks

A growing literature uses micro-level data to examine establishment-level productivity dynamics and finds substantial within-industry productivity dispersion. This paper provides an overview of a new data product, Dispersion Statistics on Productivity (DiSP), that was jointly developed and released by BLS and the Census Bureau. This new data product provides measures of productivity dispersion within narrowly defined industries by year.

Much of the paper discusses the methodology used to produce this data product. We compare our input and output measures, which we aggregated from microdata, to official BLS aggregates at the industry and manufacturing-wide level. Not surprisingly, we find some differences between BLS industry-level data and micro-aggregated ASM data. However, in general, we find high correlations between BLS and micro-aggregated outputs and inputs (for example, at the total manufacturing level, the correlation between the BLS published series and the micro-aggregated data for output and hours growth are both about 0.9).

Using measures of inputs and output, we develop measures of labor productivity (output per hour) and TFP (output per unit of combined inputs) and examine some of their properties. Correlations between BLS and micro-aggregated labor productivity growth are also reasonably high and especially high for TFP growth (e.g., at the total manufacturing level, the TFP growth correlation is 0.94).

Illustrating the properties of the new data product, we find large within-industry dispersion in labor productivity: an establishment at the 75th percentile of the productivity

³⁸ The microdata underlying the dispersion statistics has been used to explore these issues. For example, Foster et al. (2016a) and Decker et al. (2020) examined the changing relationship between productivity, survival, and growth over the cycle and in terms of secular trends.

distribution is about 2.4 times as productive as one at the 25th percentile, on average. For TFP, we find that the analogous ratio is 1.7. These patterns show enormous differences in measures of business performance across establishments in the same narrowly defined industry and year. Differences may stem from many factors, but they highlight both great potential for growth (e.g., if the gaps between high- and low-productivity businesses could be reduced) and also possible sources of frictions or distortions that are impeding a more efficient allocation of resources.

As the title of our paper suggests, we find significant dispersion in within-industry dispersion across industries. For the top quartile of industries, the ratio of TFP across establishments implied by the IQR exceeds 1.7, while for the bottom quartile of industries the ratio is lower than 1.4. Dispersion in dispersion over time is small in comparison but is still important. As has been found in previous research, we find rising dispersion in both labor productivity and TFP, and that TFP dispersion is countercyclical.

Our results also indicate that average dispersion depends on where we measure it: average dispersion is greater as we move further away from the center of the within-industry productivity distribution. Specifically, average productivity differences across establishments (especially for TFP) are largest in the right tail of the productivity distribution. Similar to what we find for average dispersion, the dynamics of these measures depend on where we measure productivity differences. We find evidence that dispersion among the most-productive establishments has been increasing during our sample period, while differences among the least-productive establishments do not show these patterns. This suggests that positively trending dispersion found in earlier studies may be a consequence of the dynamics among the most-productive establishments. The role of different establishment types is an interesting topic for future research, because assessing their relative importance would help us to better understand the drivers of productivity growth. A similarly promising area of

establishment-level productivity analysis would be to explore the role higher moments of the within-industry productivity distribution play in this regard.

Our analysis suggests that these patterns are sensitive to how dispersion is measured. We find that activity weights generally imply smaller, less volatile productivity differences among establishments for the entire distribution. We also find that, on average, activity-weighted dispersion among more-productive establishments shows a more pronounced positive trend.

Our exploratory analysis of productivity growth and dispersion indicates that dispersion and productivity growth in high-tech industries are positively correlated before the Great Recession, while this relationship reverses post-recession. While this analysis is only descriptive, it reveals systematic patterns between these key moments of the within-industry productivity distribution.

In future work, we plan to explore extending the data product in several directions. As noted, we plan to release statistics on the tails of the productivity distribution. Another area of exploration is to release statistics by additional characteristics, such as firm age and firm size. Research has shown that young businesses exhibit especially high productivity dispersion. This may reflect greater experimentation by young businesses as well as greater challenges that young businesses face in changing the scale of their operations. We also plan to use information on the occupational mix of establishments from the Occupational Employment and Wage Statistics matched to establishments in the ASM/CM to quantify the effect of labor heterogeneity on productivity dispersion.³⁹ In addition, BLS and the Census Bureau have begun work on producing dispersion statistics for retail trade industries.

³⁹ See Blackwood et al. (2022) for further discussion. An interesting finding in that research is that differences in within-industry productivity across establishments are highly correlated with differences in within-industry differences across establishments in indices of the skill and task mix of workers. This analysis highlights that novel insights can emerge from the public domain DiSP statistics while also motivating improvements in the DiSP statistics for the future.

Finally, we acknowledge that our measures of dispersion do not account for investments in intangibles, as the ASM only contains information on the book value of tangible assets and investment, and that it is possible our results could be affected if we were able to control for them. We leave the inclusion of intangibles in the measurement of productivity dispersion for future research.⁴⁰

⁴⁰ It would be interesting to explore the relationship between DiSP statistics and measures of within-industry dispersion in intangible investment similar to what Blackwood et al. (2022) did with skill and task dispersion.

References


- Asker, J., A. Collard-Wexler, and J. De Loecker, "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 122(5), 1013–1063, 2014. <https://doi.org/10.1086/677072>.
- Baily, M. N., C. Hulten, and D. Campbell, "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity*, 23 (Microeconomics), 187–267, 1992. https://www.brookings.edu/wp-content/uploads/1992/01/1992_bpeamicro_baily.pdf.
- Barth, E., A. Bryson, J. C. Davis, and R. Freeman, "It's Where You Work: Increases in Earnings Dispersion across Establishments and Individuals in the United States," *Journal of Labor Economics*, 34(S2 Part 2), S67–S97, 2016. <https://doi.org/10.1086/684045>.
- Basu, S. and J. G. Fernald, "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy*, 105(2), 249–283, 1997. <https://doi.org/10.1086/262073>.
- Becker, R. A., J. Haltiwanger, R. S. Jarmin, S. D. Klimek, and D. J. Wilson, "Micro and Macro Data Integration: The Case of Capital," in D. W. Jorgensen, J. S. Landefeld, and W. D. Nordhaus (eds), *A New Architecture for the National Accounts*, University of Chicago Press, Chicago, IL, 541–609, 2006. www.nber.org/system/files/chapters/c0146/c0146.pdf.
- Bils, M., P. J. Klenow, and C. Ruane, "Misallocation or Mismeasurement?," *Journal of Monetary Economics*, 124S, S39–S56, 2021. <https://doi.org/10.1016/j.jmoneco.2021.09.004>.
- Blackwood, G. J., L. S. Foster, C. A. Grim, J. Haltiwanger, and Z. Wolf, "Macro and Micro Dynamics of Productivity: From Devilish Details to Insights," *American Economic Journal: Macroeconomics*, 13(3), 142–172, 2021. <https://doi.org/10.1257/mac.20170282>.
- Blackwood, G. J., C. Cunningham, M. Dey, C. A. Grim, L. S. Foster, J. Haltiwanger, R. Nesbit, S. W. Pabilonia, J. Stewart, C. Tuttle, and Z. Wolf, "Opening the Black Box: Task and Skill Mix Dispersion," Paper presented at NBER/CRIW Conference on Technology, Productivity and Economic Growth, March 2022. <https://www.nber.org/conferences/criw-conference-technology-productivity-and-economic-growth-spring-2022>.
- Bloom, N., "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685, 2009. <https://doi.org/10.3982/ECTA6248>
- Burdett, K. and D. T. Mortensen, "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, 39(2), 257–273, 1998. <https://doi.org/10.2307/2527292>.
- Bureau of Labor Statistics, U.S. Department of Labor, and the Census Bureau, U.S. Department of Commerce, "Dispersion Statistics on Productivity," 1997–2016. <https://www.bls.gov/productivity/articles-and-research/dispersion-statistics-on-productivity/> and <https://www.census.gov/disp>.
- Bureau of Labor Statistics, U.S. Department of Labor, "Labor Force Statistics from the Current Population Survey: National (Unadjusted) Unemployment Rate – LNU04000000," 1996–2016. <https://data.bls.gov/timeseries/LNU04000000> (accessed March 26, 2021).

- Cao, K., K. J. Fox, and F. Soriano, "Measuring Firm-Level Capital and Productivity Using Australian Integrated Microdata," Paper presented at the Sixth World KLEMS Conference, 2021.
https://scholar.harvard.edu/files/jorgenson/files/foxcaosoriano_austrianmicrodata_klems_2021.pdf.
- Caves, D. W., L. R. Christensen, and W. E. Diewert, "Multilateral Comparisons of Output, Input, and Productivity Using Superlative Index Numbers," *The Economic Journal*, 92(365) (March 1982), 73–86, 1982. <https://doi.org/10.2307/2232257>.
- Chow, M., T. C. Fort, C. Goetz, N. Goldschlag, J. Lawrence, E. R. Perlman, M. Stinson, and T. K. White, "Redesigning the Longitudinal Business Database," Center for Economic Studies Discussion Paper, No. 21-08, 2021. <https://www2.census.gov/ces/wp/2021/CES-WP-21-08.pdf>
- Cunningham, C., L. Foster, C. Grim, J. Haltiwanger, S. W. Pablonia, J. Stewart, and Z. Wolf. 2021. "Chaos Before Order: Productivity Patterns in U.S. Manufacturing," *International Productivity Monitor*, 41, 138–152, 2021.
http://www.csls.ca/ipm/41/IPM_41_Cunningham.pdf
- Davis, S. J. and J. Haltiwanger, "Wage Dispersion Within and Between U.S. Manufacturing Plants, 1963–86," in M. Baily and C. Winston (eds.), *Brookings Papers on Economic Activity, Microeconomics*, 1991, 115–200.
- Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda, "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness," *American Economic Review*, 110(12), 3952–3990, 2020. <https://doi.org/10.1257/aer.20190680>.
- De Loecker, J., J. Eeckhout, and G. Unger, "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics*, 135(2), 561–644, 2020. <https://doi.org/10.1093/qje/qjz041>.
- Diewert, W. E. and K. J. Fox, "On Measuring the Contribution of Entering and Exiting Firms to Aggregate Productivity Growth," in W. E. Diewert, B. M. Balk, D. Fixler, K. J. Fox and A. Nakamura (eds.), *Price and Productivity Measurement - Volume 6: Index Number Theory*, Trafford Publishing, Victoria, Canada, 41–66, 2010.
- Diewert, W. E. and K. J. Fox, "A Decomposition of US Business Sector TFP growth into Technical Progress and Cost Efficiency Components," *Journal of Productivity Analysis*, 50, 71–84, 2018. <https://doi.org/10.1007/s11123-018-0535-9>.
- Dunne, T., L. Foster, J. Haltiwanger, and K. R. Troske, "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment," *Journal of Labor Economics*, 22(2), 397–430, 2004. <https://doi.org/10.1086/381255>.
- Eldridge, L. P., M. E. Manser, and P. F. Otto, "Alternative Measures of Supervisory Employee Hours and Productivity Growth," *Monthly Labor Review*, 27(4), 9–28, 2004.
www.bls.gov/opub/mlr/2004/04/art2full.pdf.
- Elvery, J., L. Foster, C. J. Krizan, and D. Talan, "Preliminary Micro Data Results from the Business List Comparison Project," Joint Statistical Meetings Proceedings, Business and Economics Statistics Section, American Statistical Association, Alexandria, VA, 2006.

- Foster, L., J. Haltiwanger, and C. J. Krizan, “Aggregate Productivity Growth: Lessons from Microeconomic Evidence,” in C. R. Hulten, E. R. Dean, and M. J. Harper (eds.), *New Developments in Productivity Analysis*, University of Chicago Press, Chicago, 63, 303–372, 2001. <https://www.nber.org/books-and-chapters/new-developments-productivity-analysis/aggregate-productivity-growth-lessons-microeconomic-evidence>.
- Foster, L., C. Grim, and J. Haltiwanger, “Reallocation in the Great Recession: Cleansing or Not?,” *Journal of Labor Economics*, 34, (S1 Part 2), S293–S331, 2016a. <https://doi.org/10.1086/682397>.
- Foster, L., C. Grim, J. Haltiwanger, and Z. Wolf, “Firm-Level Dispersion in Productivity: Is the Devil in the Details?,” *American Economic Review*, 106(5), 95–98, 2016b. <https://doi.org/10.1257/aer.p20161023>.
- , “Innovation, Productivity Growth and Productivity Dispersion,” in C. Corrado, J. Haskel, J. Miranda, and D. Sichel (eds.), *Measuring and Accounting for Innovation in the Twenty-First Century*, University of Chicago Press, Chicago, 2021a. <https://www.nber.org/books-and-chapters/measuring-and-accounting-innovation-twenty-first-century/innovation-productivity-dispersion-and-productivity-growth>.
- Foster, L., J. Haltiwanger, and C. Tuttle, “Rising Markups or Changing Technology?,” Unpublished, 2021b.
- Gort, M., and S. Klepper, “Time Paths in the Diffusion of Product Innovations,” *Economic Journal*, 92(367), 630–653, 1982.
- Griliches, Z. and H. Regev, “Firm Productivity in Israeli Industry 1979–1988,” *Journal of Econometrics*, 65(1), 175–203, 1995. [https://doi.org/10.1016/0304-4076\(94\)01601-U](https://doi.org/10.1016/0304-4076(94)01601-U).
- Haltiwanger, J. and J. Spletzer, “Between Firm Changes in Earnings Inequality: The Role of Industry and Occupation Effects,” National Bureau of Economic Research Working Paper No. 26786, 2020.
- Hsieh, C. and P. J. Klenow, “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 124(4), 1403–1448, 2009. <https://doi.org/10.1162/qjec.2009.124.4.1403>.
- Jarmin, R. S. and J. Miranda, “The Longitudinal Business Database,” Center for Economic Studies Discussion Paper No. CES-WP-02-17, 2002.
- Kehrig, M., “The Cyclicalities of Productivity Dispersion,” Social Science Research Network Working Paper, 2015. <https://ssrn.com/abstract=1854401>.
- Kovarik, T. and J. Varghese, “Intrasectoral transactions: the most important productivity statistic you’ve never heard of,” *Monthly Labor Review*, September, 2019. <https://www.bls.gov/opub/mlr/2019/article/intrasectoral-transactions-the-most-important-productivity-statistic-you-never-heard-of.htm>.
- Petrin, A., T. K. White, and J. P. Reiter, “The Impact of Plant-level Resource Reallocations and Technical Progress on U.S. Macroeconomic Growth,” *Review of Economic Dynamics*, 14(1), 3–26, 2011. <https://doi.org/10.1016/j.red.2010.09.004>.
- Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. von Wachter, “Firming Up Inequality,” *The Quarterly Journal of Economics*, 134(1), 1–50, 2019. <https://doi.org/10.1093/qje/qjy025>.


- Syverson, C., “Product Substitutability and Productivity Dispersion,” *The Review of Economics and Statistics*, 86(2), 534–550, 2004a. <https://doi.org/10.1162/003465304323031094>.
- , “Market Structure and Productivity: A Concrete Example,” *Journal of Political Economy*, 112(6): 1181–1222, 2004b. <https://doi.org/10.1086/424743>.
- , “What Determines Productivity?,” *Journal of Economic Literature*, 49(2), 326–365, 2011. <https://doi.org/10.1257/jel.49.2.326>.
- von Brasch, T., D. Iancu, and T. Skjerpen, “Productivity Dispersion and Measurement Error,” *Review of Income and Wealth*, 66(4), 985–996, 2020. <https://doi.org/10.1111/roiw.12455>.
- White, T. K., J. P. Reiter, and A. Petrin, “Imputation in U.S. Manufacturing Data and Its Implications for Productivity Dispersion,” *Review of Economics and Statistics*, 100(3), 502–509, 2018. https://doi.org/10.1162/rest_a_00678.

Figure 1. Manufacturing Employment Levels, 1997–2016




Source: “BLS Employees” is the annual average of the not-seasonally-adjusted employment in manufacturing [CEU3000000001, Current Employment Statistics program]. “ASM Published” is the published aggregate employment series from the ASM. “ASM smpl weighted” total employment is the micro-aggregated series calculated using ASM sample weights. “CMP propensity score weighted” is calculated using our estimated IPW.


Figure 2. Manufacturing Output and Input Growth Rates (in percent), 1998–2016




a) Output




b) Hours



c) Capital




d) Materials




e) Energy

Source: “BLS” is the authors’ calculations from Industry Productivity Program data. “CMP propensity score weighted” is the authors’ calculations on the ASM. “NBER” is the authors’ calculations on the NBER-CES database

Figure 3. Productivity Growth by Data Source and Measure (in percent), 1998–2016




a) Output per hour




b) Total factor productivity

Source: See the notes to Figure 2.

Figure 4. Distribution of IQR of Productivity, 1997–2016



(a) Output per hour




(b) Total factor productivity


Notes: Within-industry productivity moments are created at the four-digit NAICS level, weighted using IPW. Annual descriptive statistics of industry dispersion are unweighted.

Source: Authors' calculations on the ASM.

Figure 5. Distribution of *Activity-Weighted* IQR of Productivity, 1997–2016




(a) Output per hour (hours-weighted)




(b) Total factor productivity (composite-input-weighted)

Source: See notes to Figure 4.

Figure 6. Distribution of 99–90 Difference of Productivity, 1997–2016




(a) Output per hour




(b) Total factor productivity

Source: See notes to Figure 4.

Figure 7. Distribution of *Activity-Weighted* 99–90 Difference of Productivity, 1997–2016




(a) Output per hour (hours-weighted)




(b) Total factor productivity (composite-input-weighted)

Source: See notes to Figure 4.

Figure 8. Distribution of 10–1 Difference of Productivity, 1997–2016




(a) Output per hour




(b) Total factor productivity

Source: See notes to Figure 4.

Figure 9. Distribution of *Activity-Weighted* 10–1 Difference of Productivity, 1997–2016



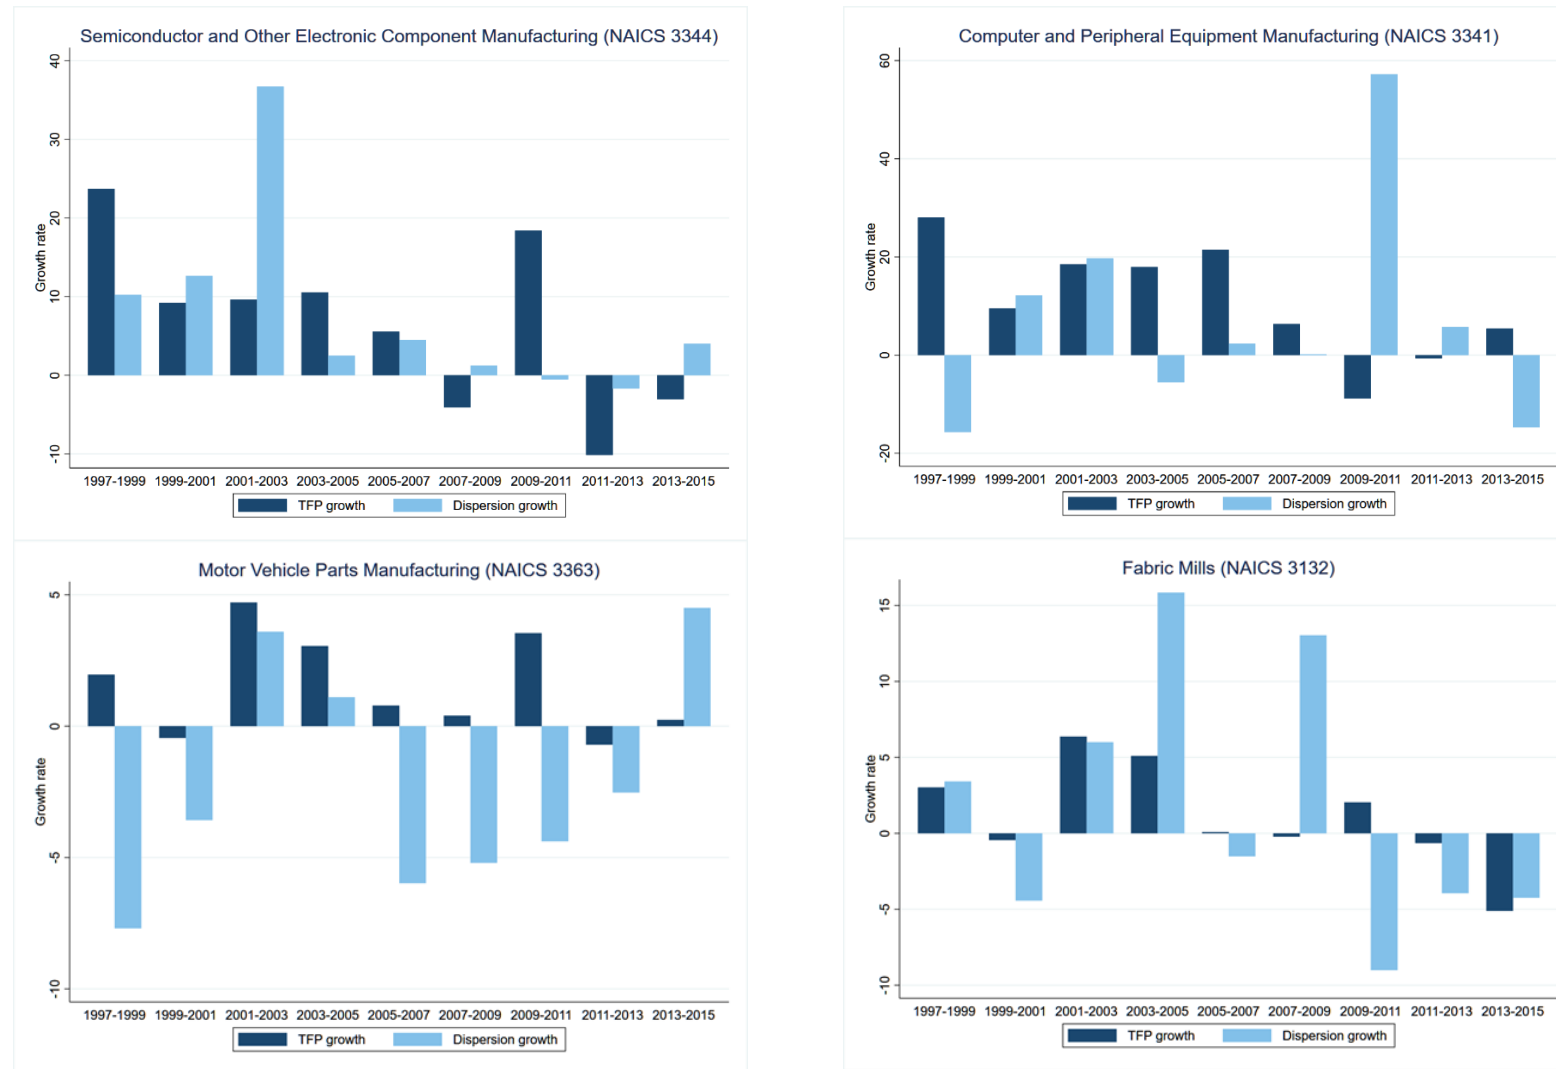
(a) Output per hour (hours-weighted)



(b) Total factor productivity (composite-input-weighted)

Source: See notes to Figure 4.

Figure 10. Growth in TFP and Dispersion in Select Industries (1997–2015)



Source: Author's calculations on the DiSP data and BLS Industry Productivity Statistics

Table 1. Input and Output Correlations between BLS, CMP, and NBER (1997–2016)

	BLS/CMP	BLS/NBER	CMP/NBER
Total Manufacturing			
Hours worked, levels	0.986	0.994	0.996
Hours worked, growth	0.930	0.909	0.923
Capital, levels	-0.184	0.880	0.328
Capital, growth	0.330	0.643	0.585
Energy, levels	0.977	0.997	0.978
Energy, growth	0.958	0.985	0.948
Materials, levels	0.931	0.996	0.963
Materials, growth	0.858	0.992	0.880
Output, levels	0.945	0.996	0.961
Output, growth	0.894	0.993	0.926
Average of Four-digit NAICS			
Hours worked, levels	0.803	0.889	0.883
Hours worked, growth	0.457	0.632	0.599
Capital, levels	0.521	0.714	0.489
Capital, growth	0.265	0.565	0.263
Energy, levels	0.861	0.986	0.841
Energy, growth	0.714	0.714	0.709
Materials, levels	0.837	0.962	0.843
Materials, growth	0.659	0.937	0.661
Output, levels	0.838	0.985	0.850
Output, growth	0.676	0.951	0.675

Source: Authors' calculations on the ASM.

Table 2. Productivity Growth Correlations between BLS, CMP, and NBER (1997–2016)

	BLS/CMP	BLS/NBER	CMP/NBER
Labor productivity (Total Manufacturing)	0.705	0.818	0.735
Labor productivity (Average of Four-digit NAICS)	0.465	0.619	0.660
Total factor productivity (Total Manufacturing)	0.935	0.991	0.960
Total factor productivity (Average of Four-digit NAICS)	0.786	0.896	0.810

Source: Authors' calculations on the ASM.

Table 3. Summary of Within-Industry Productivity Distributions (1997–2016)

Within-Industry Productivity Moment	Mean	Standard Deviation	IQR
Labor Productivity			
IQR	0.898	0.290	0.322
90-10 differential	1.773	0.476	0.613
Standard deviation	0.684	0.167	0.222
99-90 differential	0.732	0.279	0.333
10-1 differential	0.550	0.267	0.275
Total Factor Productivity			
IQR	0.520	0.222	0.205
90-10 differential	1.078	0.393	0.371
Standard deviation	0.460	0.152	0.161
99-90 differential	0.866	0.512	0.546
10-1 differential	0.301	0.181	0.153

Notes: Log labor productivity (LP) is calculated as log (output/hours) where hours are BLS-adjusted total hours. The four-digit NAICS industry mean log LP is subtracted off establishment-level log LP. Within-industry productivity moments are created at the four-digit NAICS level using IPW. Annual summary statistics of these industry statistics are then created weighting each industry equally. The numbers shown are means of the annual summary statistic values from 1997 to 2016, weighting each year equally.

Source: Authors' calculations on the ASM.

Table 4. Probability of Transition across Quintiles of the Cross-Industry Distribution of Dispersion (annual averages between 1997 and 2016)

IQR						IQR (Activity weighted)					
Output per Hour											
	1	2	3	4	5		1	2	3	4	5
1	0.63	0.24	0.08	0.03	0.01	1	0.70	0.22	0.06	0.01	0.00
2	0.23	0.41	0.22	0.10	0.02	2	0.22	0.49	0.24	0.04	0.00
3	0.10	0.25	0.41	0.25	0.03	3	0.06	0.25	0.56	0.15	0.01
4	0.03	0.08	0.25	0.44	0.19	4	0.01	0.04	0.14	0.65	0.15
5	0.01	0.03	0.03	0.18	0.74	5	0.00	0.00	0.01	0.15	0.84

Total Factor Productivity											
	1	2	3	4	5		1	2	3	4	5
1	0.64	0.23	0.07	0.04	0.01	1	0.75	0.19	0.03	0.01	0.01
2	0.22	0.42	0.28	0.05	0.02	2	0.19	0.53	0.21	0.05	0.00
3	0.08	0.26	0.40	0.26	0.04	3	0.03	0.24	0.53	0.20	0.03
4	0.04	0.08	0.20	0.47	0.18	4	0.02	0.02	0.20	0.56	0.17
5	0.01	0.02	0.05	0.16	0.75	5	0.01	0.01	0.03	0.17	0.79

Notes: Rows index quintiles in $t-1$, columns index quintiles in t . Probabilities in each table are normalized by column sums, i.e., column elements sum to one, apart from rounding.

Source: Authors' calculations on the ASM.

Table 5. Common Secular Trends and Cyclicalities in IQR Dispersion

	Dependent Variable			
	Labor Productivity		Total Factor Productivity	
	Non-activity weighted	Activity Weighted	Non-activity weighted	Activity Weighted
Change in unemployment rate (in 100s)	-1.182 (0.390)	-0.093 (0.293)	1.061 (0.458)	0.686 (0.338)
Trend	0.004 (0.001)	0.004 (0.001)	0.004 (0.001)	0.002 (0.002)
N	1,720	1,720	1,720	1,720
R-squared	0.647	0.809	0.644	0.761

Notes: Sample is the 4-digit IQR dispersion statistics from 1997–2016. All specifications control for four-digit NAICS industry effects. Robust standard errors in parentheses are clustered at the industry level.

Source: Authors' calculations on the DiSP data and BLS unemployment rate.

Table 6. Relationship between Productivity and Establishment Characteristics, 1997–2016

Characteristic	R ²
	LogLP (demeaned)
State	0.003
Size class	0.014
Age class	0.007
	LogTFP (demeaned)
State	0.006
Size class	0.004
Age class	0.001

Notes: LogLP (demeaned): log labor productivity, industry and year effects are removed, LogTFP (demeaned): log TFP, industry and year effects are removed, naics4: four-digit NAICS code, year: time identifier, state: Federal Information Processing Standard state code, size class: employment size class with thresholds 20, 50, 100, 250, and 500, age class: establishment age class with thresholds 1, 2, 3, 4, 5, 10, and 15. Each row shows the R² from a separate regression; in each instance, the p-value is less than 0.0001.

Source: Authors' calculations on the ASM.

Table 7. BLS Productivity Growth and Dispersion Growth Before and After the Great Recession

	Pre-Recession Period (1999–2007)				Post-Recession Period (2009–2015)			
	(1) LP	(2) LP	(3) TFP	(4) TFP	(5) LP	(6) LP	(7) TFP	(8) TFP
Dispersion	0.188 (0.073)	0.110 (0.069)	0.062 (0.021)	0.043 (0.022)	0.116 (0.057)	0.150 (0.065)	-0.038 (0.027)	-0.001 (0.017)
Lagged Dispersion	0.012 (0.064)	-0.082 (0.046)	0.005 (0.022)	-0.016 (0.015)	0.126 (0.072)	0.143 (0.076)	-0.027 (0.012)	-0.016 (0.014)
Tech x Dispersion		0.317 (0.110)		0.054 (0.051)		-0.117 (0.115)		-0.123 (0.041)
Tech x Lagged Dispersion		0.509 (0.093)		0.060 (0.062)		-0.081 (0.155)		-0.049 (0.027)
Joint Hypothesis Tests:								
[Dispersion + Tech x Dispersion]		0.427 (0.088)		0.097 (0.047)		0.033 (0.095)		-0.124 (0.038)
[Lagged Dispersion + Tech x Lagged Dispersion]		0.427 (0.078)		0.044 (0.061)		0.062 (0.137)		-0.064 (0.024)
<i>N</i>	344	344	344	344	258	258	258	258
<i>R</i> -squared	0.141	0.227	0.118	0.154	0.155	0.166	0.178	0.211

Notes: In these regressions, average annual industry-level productivity growth rates for two-year subperiods are regressed on average annual industry-level dispersion growth rates for two-year subperiods. Dispersion is defined as the activity-weighted IQR for an industry in a year. Controls include a constant and period effects. Robust standard errors in parentheses are clustered at the industry level.

Source: Author's calculations on the DiSP data and BLS Industry Productivity Statistics.

Appendix A

Table A1. Table of Acronyms

Acronym	Meaning
ACES	Annual Capital Expenditures Survey
ASM	Annual Survey of Manufactures
AWH	Average weekly hours
BEA	Bureau of Economic Analysis
BLS	Bureau of Labor Statistics
CES	Current Employment Statistics
CM	Census of Manufactures
CMP	Collaborative Micro-productivity Project
CPS	Current Population Survey
CR	Cost of resales
DF	Changes in finished goods
DiSP	Dispersion Statistics on Productivity
DW	Changes in work-in-process inventories
E	Cost of electricity and cost of fuels
FIB	Beginning-of-year finished goods inventory
FIE	End-of-year finished goods inventory
FSRDC	Federal Statistical Research Data Center
GDP	Gross domestic product
IPW	Inverse propensity score weights
IQR	Interquartile range
LBD	Longitudinal Business Database
LP	Labor productivity
M	Cost of materials, cost of resales, and cost of contract work
NAICS	North American Industry Classification System
NBER-CES	National Bureau of Economic Research and Center for Economic Studies
NCS	National Compensation Survey
PH	Production worker hours
PISHIP	Deflator for value of shipments
PW	Number of production workers
Q	Output
TE	Total employment
TFP	Total factor productivity
TH	Total hours worked
TVS	Total value of shipments
WIB	Beginning-of-year work-in-process inventories
WIE	End-of-year work-in-process inventories

Table A2. Summary of Variables Used in Selected Tables and Figures

	Table 1 Correlations	Table 2 Correlations	Table 3 Dispersion	Figure 1 Employment	Figure 2 Hours	Figure 2 Output	Figure 3 Productivity
				Comparisons			
BLS implicit price deflator used for all estimates (except for capital)	yes	yes				yes	yes
Shipments deflator used to deflate output			yes				
Cost of resales (CR) removed from CMP	yes	yes	yes			yes	yes
Employees only			yes	yes		N/A	
Include BLS employees and self-employed (SE) and unpaid family workers (UFW) in BLS data only	yes	yes			yes	N/A	yes
CPS nonproduction/production hours ratio (even for NBER hours)	yes	yes	yes		yes		yes
BLS intrasectorals included	yes	yes				yes	yes

Appendix B

B.1. Properties of ASM samples

The ASM is a five-year panel of roughly 50,000–70,000 manufacturing establishments. It is a sample of establishments drawn from the manufacturing portion of the Census Bureau’s Business Register using a probability proportional to size sampling scheme.⁴¹ The largest establishments are sampled with certainty and are included in every panel.⁴² Smaller establishments are sampled with a probability less than one, where the probability increases with establishment size (measured by the value of shipments). The smallest single-unit establishments, which are part of the “non-mail” stratum, are not mailed a form but they are included in the estimates. The Census Bureau uses administrative records for payroll, employment, industry, and location from the administrative data for the smallest single-unit establishments, while total value of shipments is imputed using industry averages.⁴³

The ASM sample is refreshed every five years. New ASM panels are drawn from the Economic Census and begin 2 years after the Census from which it was drawn (years ending in “4” and “9”). The sample is also updated annually to include new establishments that are identified on the Census Bureau’s Business Register. The Business Register is updated with information from the Economic Census as well as administrative records from the IRS and the Census Bureau’s annual Report of Organization (formerly called the Company Organization Survey).

⁴¹ For more information about the ASM, see <http://www.census.gov/manufacturing/asm/>.

⁴² Prior to 1999, certainty units were establishments with 250 or more employees. In 1999, the cutoff was increased to 500 employees and, in 2004, it was increased again to 1,000 employees. Currently, the 10 largest establishments in an industry are also sampled with certainty. In addition to establishment size, certainty criteria include other characteristics such as industry, cell size, or energy use. For example, computers, flat-glass, sugar, and small industries (with less than 20 establishments), or establishments with large inventories, assets, fuel/electric expenditures are also sampled with certainty.

⁴³ Non-mail cases are included in the official estimates and have a weight of one. The survey is designed to tabulate cases from the mail and the non-mail component. The mail component was not designed to estimate the total population.

Data for the ASM are collected in all years except for years ending in “2” and “7”, when the ASM data are collected as part of the Economic Census. Data on payroll, employment, industry, and geography for establishments in the non-mail stratum are obtained from administrative records.⁴⁴

The ASM sample is designed to estimate unbiased national level estimates of a skewed population. For example, in the 2014 ASM panel, large establishments sampled with certainty account for approximately 72% of the total value of shipments in the 2012 CM; non-certainty establishments are sampled with probabilities from 0.05 to 1.00.⁴⁵ This sample design implies that the establishment counts in various size bins may not reflect those calculated from the LBD.

The ASM sample weights, which are inversely proportional to a shipments-based establishment size measure, could in principle be used to correct for the effects of the ASM sample design. However, the sample design implies that the weighted sum of shipments from the mail stratum only will not match published totals.⁴⁶

Another important aspect of the sample design is that the composition of establishments changes over time and between sample selections. Any weighting procedure aiming at creating unbiased estimates of productivity dispersion should account for the fact that the sampling probabilities, and therefore the composition of the ASM, change every five years. In addition, sampling and non-mail stratum thresholds vary across years.

B.2. Establishment Characteristics and the Probability of Selection into the ASM

The ASM’s sample design has important implications for our analysis. For example, the sum of the ASM sample weighted employment or sales might equal total employment or total

⁴⁴ Federal regulations require the Census Bureau to limit survey response burden.

⁴⁵ Source - ASM Methodology website, <https://www.census.gov/programs-surveys/asm/technical-documentation/methodology.html> (accessed September 16, 2020).

⁴⁶ As mentioned above, only the mail component together with the adjustment for the non-mail stratum yields unbiased estimates of the total population. See Davis et al. (1996) for more details.

sales. However, it is not clear that the ASM sample weights are appropriate for our analysis. This section is devoted to describing our weighting procedure.

To address the effects of the ASM's sample design, we construct IPW using the LBD. Propensity scores are estimated from a logistic regression in which we model the relationship between establishment characteristics and the probability that an establishment is selected into the ASM. We start by matching establishments in the ASM to LBD establishments by year and "LBD Number."⁴⁷ Our dependent variable is a dummy variable that equals one if the establishment is in both the ASM and the LBD for that year and zero if the establishment is only in the LBD. For establishments in the non-mail stratum, the dummy variable is set to zero.

The set of regressors includes dummy variables that classify each establishment based on its employment and payroll size class, whether the establishment is part of a multi-unit entity, the establishment's industry code, and the interaction between industry and employment size effects. Including industry-size interactions allows us to estimate industry-specific size distributions. These variables are obvious candidates for our logistic regressions because the probability of selection into the ASM sample and the cutoff for the non-mail stratum in the ASM vary by industry and size.

We define industries at the three-digit NAICS level because the interaction of size indicators and more narrowly defined industry codes leads to empty cells in smaller industries. Empty-size bins imply that the size distribution cannot be estimated in these industries.⁴⁸ When the size distribution cannot be estimated for an industry, propensity scores cannot be calculated because maximum-likelihood estimates of the size effects do not exist. Empty cells can, in principle, be avoided by collapsing size bins, combining similar narrowly defined industries, or allowing bin definitions to vary across industries. We experimented with the number and

⁴⁷ The LBD Number is an establishment identifier that is consistently defined across both datasets. Although linking the datasets by LBD Number is straightforward, a small percentage of establishment-year observations do not match due to timing issues between the ASM and the LBD.

⁴⁸ The size distribution cannot be estimated if all establishments are in the same size bin.

definition of the size bins and the level of industry aggregation and found that using three-digit NAICS industry codes together with four size bins allows us to estimate the size distribution in every industry and year. Allowing for more heterogeneity by using either industry-specific size bins or more narrowly defined industries leads to feasibility problems with the logistic regression.

We define the size bins so that the resulting distribution allows the lowest size bins to vary over time. That is, in every year and every industry, the 50th percentile of establishments with fewer than 50 employees is used to define bins one and two. For larger establishments, the following bins are defined: 50–99, 100–199, and 200+. ⁴⁹ There are 21 three-digit NAICS industries in the 2002 classification system, which results in 105 industry-specific size distributions. We include a continuous size measure to allow the weights to vary within these cells. This is necessary to account for possible within-cell compositional changes. Adding five payroll classes and two groups related to multi-unit status increases the number of cells to 113. ⁵⁰

The 2002 change in the industry classification system resulted in missing NAICS-2002 codes for a nontrivial number of establishments in the LBD between 1997 and 2001. For example, the NAICS code is missing if an establishment exited prior to 2002. For these observations, we used imputed NAICS codes. ⁵¹ From 2002 on, NAICS codes are available for all establishments in the LBD.

B.3 Comparison of Hours Measures

In this study, we use hours data from ASM, augmented with data from the CPS. However, for official estimates of productivity growth, BLS uses the CES as its primary source of hours data. Although the CES and ASM are establishment surveys, the two surveys differ in

⁴⁹ The payroll size classes are 0–200, 201–500, 501–1000, 1001–5000, and 5001+.

⁵⁰ If we were to use 4-digit industry, the number of cells would increase significantly. There are 86 4-digit NAICS industries implying 86 different size distributions and 430 industry-size cells. Such an increase in the number of cells yields empty size bins in several industries.

⁵¹ NAICS codes are imputed using a method described in Fort (2013).

what hours data they collect and how they collect it. The best information on these differences comes from studies completed in the early 2000s (Goldenberg and Willimack, 2003; Fisher et al., 2001). These studies do a nice job of summarizing the differences between the two surveys and how those differences affect estimates of hours worked. In this appendix, we summarize that research and discuss the implications for comparing our estimates to published BLS estimates.

There are some general differences between the two surveys that are worth noting. First, the ASM is an annual survey, whereas the CES is conducted monthly. As a result, the reference periods of the two surveys differ. The reference period for the CES is the pay period that includes the 12th of the month. The CES collects data on the total number of employees, hours for all employees since 2006, the number of production workers, production worker payroll, and production worker hours.

In contrast, the ASM has different reference periods for different data elements. For production worker employment, the ASM reference period is the pay period that includes the 12th of the month in the months of March, May, August, and November. These quarterly reports are then averaged into an annual number. The ASM collects employment data for other employees only for the pay period that includes March 12th. The implicit assumption is that nonproduction worker employment does not vary much over the year. Total employment is not collected directly, but rather is equal to the sum of the total number of nonproduction workers in March and the annual average of quarterly production worker employment.

Annual total employment in the two surveys can differ if there are seasonal patterns in production worker employment that are missed in the ASM's quarterly reports or if there is a seasonal pattern to nonproduction worker employment. We examined this issue using monthly CES data. Specifically, we calculated the average employment for each quarter using CES data, and then calculated the ratio of average quarterly employment to CES employment in the ASM

reference months (March, May, August, and November). The ratios were very close to one, showing that estimates of average annual employment are the same whether we use four quarterly reports or 12 monthly reports.

There are greater differences in the hours data collected in the two surveys. There are two possible reasons for this. First, the two surveys use different concepts. The ASM asks employers to report hours *worked*, whereas the CES collects hours *paid*. The main difference is that the CES hours data include holidays, annual leave, and sick leave that were paid but not worked. Thus, we would expect total annual hours reported in the CES to exceed total annual hours in the ASM. For productivity measurement, hours worked is the correct concept, which is why BLS adjusts the CES data using hours-worked-to-hours-paid ratios from the NCS.

Second, the two surveys differ in how they ask respondents to report hours. The ASM asks respondents to report total annual production worker hours. The CES asks respondents to report total employment and hours for production workers for the pay period that includes the 12th of the month. The hours reports are converted to a weekly number using conversion factors that vary with the number of workdays in the month. Apart from the difference in concept, these two approaches to collecting hours data could result in different estimates of total annual hours. Research by Frazis and Stewart (2004) has shown that people work longer hours during the week that includes the 12th of the month.⁵² This would also lead to annual hours in the CES being higher than in the ASM. Neither survey collects hours data for nonproduction workers.⁵³ As noted in the text, nonproduction worker hours are estimated using data from the CPS.

⁵² Their research examined the accuracy of CPS hours reports by comparing the CPS hours data to hours data from the American Time Use Survey (ATUS). They found that reports were on average consistent, but that some groups overreported their hours (college educated) and others underreported their hours (less than college). Research by Eldridge et al. (2022) found differences in hours reports for production and nonproduction workers that are consistent with the findings of Frazis and Stewart (2004).

⁵³ The CES began collecting all employee hours in 2006.

Appendix References

- Davis, S. J., J. C. Haltiwanger, and S. Schuh, *Job Creation and Destruction*, Cambridge, M.A.: MIT Press, 1996.
- Eldridge, L. P., S. W. Pabilonia, and J. Stewart, “Improving Estimates of Hours Worked for U.S. Productivity Measurement,” *Monthly Labor Review*, forthcoming, 2022.
- Fisher, S., K. Goldenberg, E. O’Brien, C. Tucker, and D. Willimack, “Measuring Employee Hours in Government Surveys,” Paper presented at the Federal Economic Statistics Advisory Committee (FESAC), Washington, DC, 2001.
- Fort, T., “NAICS 2002 Code Assignments to the Longitudinal Business Database,” Unpublished, 2013.
- Frazis, H. and J. Stewart, “What Can Time-use Data Tell Us about Hours of Work?,” *Monthly Labor Review*, 127(12), 3–9, 2004. <https://www.bls.gov/opub/mlr/2004/12/art1full.pdf>.
- Goldenberg, K. L., and D. K. Willimack, “Measurement Differences in Key Economic Indicators,” Paper presented at the Joint Statistical Meetings of the American Statistical Association, San Francisco, CA, 2003.